Bitcoin & Quantum Resistance: State of the Field

A Technical Assessment — March 2026

Research report by Richard DeVaul & GLaDOS First published: March 5, 2026 Last updated: March 5, 2026

Executive Summary

Quantum computing poses a real but non-imminent threat to Bitcoin's cryptographic foundations. The primary vulnerability is ECDSA (transaction signatures), not SHA-256 (mining). No cryptographically relevant quantum computer (CRQC) exists today, and expert consensus places the earliest plausible timeline at 2030–2035. However, given Bitcoin's multi-year upgrade cycles, the community must begin serious protocol work now. BIP 360 ("QuBit") is the most concrete proposal to date, merged into the BIP repository in February 2026. The politics are contentious — particularly around what to do with ~6.5 million BTC in quantum-vulnerable addresses, including Satoshi's estimated 1.1 million BTC.

1. The Quantum Threat to Bitcoin — What Is Actually at Risk

Bitcoin's cryptographic security rests on two distinct pillars, each facing a different quantum threat profile: **elliptic curve digital signature algorithm (ECDSA)** used for transaction authorization, and **SHA-256** used in proof-of-work mining and Merkle tree construction.

1.1 ECDSA and Shor's Algorithm

Bitcoin's ECDSA uses the secp256k1 curve, relying on the intractability of the elliptic curve discrete logarithm problem (ECDLP). A sufficiently powerful quantum computer running **Shor's algorithm** can solve ECDLP in polynomial time — essentially deriving a private key from a public key. This is the existential threat.

The attack surface has two tiers:

Exposed public keys (offline attack): Any address from which Bitcoin has been *sent* reveals its public key in the spending transaction. This includes all legacy P2PK (Pay-to-Public-Key) addresses — notably the Satoshi-era coinbase outputs — plus any reused P2PKH or Bech32 address. The Human Rights Foundation estimated in a 2025 report that approximately **6.51 million BTC** sits in quantum-vulnerable addresses; of that, ~1.72 million BTC (including Satoshi Nakamoto's estimated 1.1 million BTC) is in unmoved, likely-lost coins with permanently exposed public keys.

In-flight transactions (online attack): During the brief window between broadcast and confirmation, an ECDSA signature is visible in the mempool. An attacker with a fast enough quantum computer could, in theory, derive the private key in that window and submit a competing double-spend. This requires quantum computation within the ~10-minute block interval — a much higher bar than the offline attack on dormant exposed keys.

Scale required: Breaking secp256k1 with Shor's algorithm requires an estimated **2,330 logical qubits** for the core computation. But error correction overhead is severe: estimates range from **900,000 to several million physical qubits** depending on qubit fidelity and the error-correction code used. A 2025 Google paper revised earlier estimates to ~900,000 physical qubits. A 2026 preprint ("The Pinnacle Architecture") suggested optimized architectures might reduce this below 100,000 physical qubits — though this remains unverified and is likely optimistic.

1.2 SHA-256 and Grover's Algorithm

Grover's algorithm provides a quadratic speedup for unstructured search, reducing SHA-256's effective security from 256 bits to 128 bits. This affects proof-of-work mining: a quantum miner could find hashes roughly twice as fast as a classical miner of equivalent cost. This is a *competitive* advantage, not an existential break. Bitcoin's difficulty adjustment would compensate. The consensus among cryptographers is that SHA-256 remains quantum-safe for practical purposes — doubling key/hash lengths (a soft upgrade) would fully restore original security margins if needed.

Similarly, Bitcoin's use of SHA-256d (double-SHA-256) in Merkle trees and TXID computation is not meaningfully threatened by Grover's — 128-bit post-quantum security is generally considered acceptable for these use cases.

Bottom line on threat asymmetry: The ECDSA/Shor threat is qualitatively different from the SHA-256/Grover threat. The former enables *theft*; the latter at worst enables *faster mining*. All serious discussion of Bitcoin's quantum vulnerability focuses on ECDSA.

1.3 Realistic Timeline

Honest assessment of the timeline requires separating marketing from physics. As of early 2026:

System	Physical Qubits	Status
Google Willow	105	Meaningful error correction progress, nowhere near cryptographic scale
Microsoft Majorana 1	Early-stage	Topological architecture with theoretical fidelity advantages
IBM Nighthawk	~120	Incremental improvements
Required for ECDSA break	**~900,000+**	**~8,500× current state of the art**

The authoritative analysis from a16z Crypto (December 2025) concludes: *"A cryptographically relevant quantum computer (CRQC) in the 2020s is highly unlikely… The gap between demonstrating that quantum error correction works in principle, and achieving the scale needed for cryptanalysis, remains vast."*

However, 80 experts at the **July 2025 Presidio Bitcoin Quantum Summit** assessed that a CRQC capable of breaking Bitcoin's ECDSA could emerge within **5–10 years** — a 2030–2035 window. This aligns with cautious mainstream projections.

The operationally important implication: given the time required to design, test, and deploy a protocol-level cryptographic upgrade to Bitcoin (historically 2–4+ years from proposal to activation), **work must begin well before a CRQC exists.**

2. Current Efforts Underway

2.1 BIP 360 — QuBit (Pay-to-Quantum-Resistant-Hash)

The most significant concrete proposal to date is **BIP 360**, nicknamed "QuBit," authored by Hunter Beast. It was formally merged into the Bitcoin BIP GitHub repository in **February 2026** — a notable milestone, though merging a BIP does not imply adoption or even consensus, only that the proposal meets basic formatting and clarity standards for discussion.

BIP 360 proposes a new output type tentatively called **Pay-to-Merkle-Root (P2MR)**. The design philosophy mirrors Taproot's structure — a Merkle tree of spending conditions — but replaces the secp256k1 internal key with a quantum-resistant commitment. Key properties:

- **Public keys are never exposed until spending**, protecting unspent outputs
- **Backward-compatible as a soft fork** (new output type with WITNESS_VERSION)
- **Designed for gradual migration** — users opt in by moving funds to new address type
- **Supports FALCON-1024** as the primary signature scheme, with accommodations for SPHINCS+ and ML-Kyber-based constructions

2.2 QRAMP — Quantum-Resistant Address Migration Protocol

Proposed in 2025, **QRAMP** is a more aggressive companion proposal addressing the *existing* exposed-key problem. It would establish a **mandatory migration deadline**: after a defined sunset block height, coins in quantum-vulnerable address formats (P2PK, reused P2PKH) would be unspendable or would require spending to a quantum-resistant address.

This is extraordinarily controversial (see Section 4).

2.3 BIP-QShield

Drafted in July 2025, **BIP-QShield** takes the most radical position: permanently *exclude* coins in quantum-vulnerable addresses from on-chain transactions. This would effectively **freeze Satoshi's coins** and ~6.5 million other exposed BTC. It remains highly contentious and has not received meaningful technical traction.

2.4 Bitcoin Dev Mailing List Activity

2025 saw a significant uptick in quantum-related discussion on the Bitcoin-dev mailing list:

- **Jesse Posner** (Coinbase cryptographer) published research showing that HD wallets, silent payments, key aggregation, and threshold signature schemes could be compatible with some quantum-resistant signature algorithms
- **Jameson Lopp** (CasaHODL) proposed a concrete sequence of soft forks deployable *before* a CRQC arrives, gradually tightening quantum-vulnerable spend rules
- **Augustin Cruz** proposed a BIP to proactively destroy quantum-vulnerable coins
- **BTQ Technologies** (October 2025) announced a working implementation of Bitcoin using NIST-standardized ML-DSA (formerly CRYSTALS-Dilithium, now FIPS 204) replacing ECDSA, demonstrating feasibility at the node level

2.5 The Presidio Bitcoin Quantum Summit

The **July 2025 Presidio Bitcoin Quantum Summit** brought together 80 researchers and practitioners in the first major dedicated forum on the topic. The Human Rights Foundation funded original research quantifying the exposure of ~6.51 million BTC. A ScienceDirect survey published November 2025 provides the most comprehensive academic treatment to date.

3. Candidate Post-Quantum Cryptographic Schemes

NIST completed its first post-quantum standardization round in 2024, producing three signature standards relevant to Bitcoin:

Algorithm	NIST Standard	Type	Signature Size	Public Key Size	Notes
ML-DSA (Dilithium)	FIPS 204	Lattice (Module-LWE)	2,420–4,595 bytes	1,312–2,592 bytes	Simplest to implement correctly
FN-DSA (FALCON)	FIPS 206	Lattice (NTRU)	666–1,280 bytes	897–1,793 bytes	Best size profile; complex signing
SLH-DSA (SPHINCS+)	FIPS 205	Hash-based	7,856–49,856 bytes	32–64 bytes	Conservative (hash-only security)

Compare with Bitcoin's current ECDSA: ~71-byte signature, 33-byte compressed public key.

3.1 Tradeoffs for Bitcoin

FALCON-1024 (FN-DSA) is the leading candidate for Bitcoin integration, specifically cited in BIP 360. At ~1,280-byte signatures with 1,793-byte public keys, it has the best size profile among lattice schemes. The catch: FALCON requires Gaussian sampling during signing, which is notoriously difficult to implement in constant time without subtle timing side-channels. This implementation complexity is a serious concern for a security-critical application.

ML-DSA (Dilithium) is simpler to implement correctly and is the primary NIST recommendation for general use, but its ~2,420-byte signatures would roughly **30× the per-transaction signature data**. At current Bitcoin transaction throughput, this translates directly to blockchain bloat — a soft-fork-only upgrade would need to accommodate this within existing block weight limits, severely constraining transaction throughput.

SPHINCS+ (SLH-DSA) has the advantage of being purely hash-based — its security relies only on the collision resistance of SHA-256 or similar, making it conservative against any future lattice cryptanalysis breakthroughs. However, at 7,856–49,856 bytes per signature, it's effectively unusable for routine Bitcoin transactions without dramatic block size increases.

ML-KEM (Kyber) — the NIST KEM standard (FIPS 203) — is not directly applicable to Bitcoin's signature use case but is relevant for off-chain channel establishment (Lightning Network) and encrypted wallet communications.

3.2 Hybrid Approaches

A September 2025 preprint ("Hybrid Post-Quantum Signatures for Bitcoin and Ethereum") proposes signing transactions with *both* ECDSA and a PQC scheme during a transitional period. This provides security against either classical *or* quantum attackers, at the cost of double the signature overhead. The paper notes this imposes "immediate, severe costs with no tangible benefits" until a CRQC actually exists — framing it as a "defensive downgrade."

3.3 Block Weight Implications

FALCON-1024 signatures are ~18× larger than ECDSA signatures. In a standard 1-input, 1-output transaction, this balloons the transaction from ~250 bytes to roughly 3,500+ bytes. Bitcoin's current 4 million weight-unit block limit would mean roughly **1/18th the transaction throughput** for PQC-only blocks.

Accommodating PQC signatures likely requires either: 1. A **block weight increase** (contentious — echoes of the block size wars) 2. **Signature aggregation** schemes that amortize PQC signature size across multiple inputs 3. **Off-chain solutions** (Lightning) handling most transactions, with on-chain serving only as settlement layer

4. Bitcoin Politics: The Hardest Part

Technical solutions exist. The politics are where this gets difficult.

4.1 The Property Rights Question

QRAMP and BIP-QShield propose, in different ways, restricting or freezing coins in quantum-vulnerable addresses. This raises a fundamental question that has no technical answer:

Does the Bitcoin network have the right to render someone's coins unspendable?

Arguments for migration deadlines: - Without action, a quantum attacker could steal millions of BTC, crashing the network's value for everyone - A mandatory migration protects the *ecosystem*, not just individual holders - Coins in lost wallets (including Satoshi's) are already effectively unspendable; formally freezing them changes nothing practically

Arguments against: - Bitcoin's core value proposition is **censorship-resistant, permissionless money**. Freezing anyone's coins — for any reason — violates this principle. - Some holders may be unable to migrate (lost keys to multisig setups, legal holds, deceased holders' estates) - Setting a precedent for "the network can freeze your coins" opens the door to future politically-motivated freezes - **Jameson Lopp's position**: Better to let a quantum attacker claim Satoshi's coins than to set the precedent that the network can freeze coins

4.2 Soft Fork vs. Hard Fork

BIP 360 is designed as a **soft fork** — backward-compatible, new output type, voluntary migration. This is the path of least political resistance and mirrors the Taproot playbook.

A hard fork would be required for: - Mandatory migration (QRAMP at its most aggressive) - Block weight increases to accommodate PQC signature sizes - Changes to the mining algorithm (replacing SHA-256)

Given Bitcoin's governance culture, **a hard fork for quantum resistance is extremely unlikely** absent an imminent, demonstrated threat. The community will almost certainly pursue soft fork approaches until forced to do otherwise.

4.3 The Urgency Debate

"Act now" camp: The upgrade cycle for Bitcoin consensus changes is 2–4+ years from proposal to activation. If a CRQC arrives in 2030, protocol work should have been done by 2026–2028. BIP 360's merge is a step, but activation requires years of review, testing, and signaling. Every year of delay compresses the timeline.

"Don't rush" camp: Premature standardization risks picking the wrong PQC algorithm. NIST's standards are new (2024); lattice-based schemes haven't been battle-tested at scale. A hasty upgrade could introduce new vulnerabilities (e.g., FALCON's timing side-channels) worse than the theoretical quantum threat. Bitcoin's conservatism is a feature, not a bug.

"It's overhyped" camp: Current quantum hardware is ~8,500× away from cryptographic relevance. Moore's Law doesn't apply to qubit scaling (the challenges are physical, not lithographic). A 2035 CRQC is possible; a 2030 CRQC is marketing. There's time to get this right.

4.4 Current Consensus (or Lack Thereof)

There is no consensus. The community is in the "awareness and early research" phase: - BIP 360 is merged but has not entered the activation discussion - No miner signaling, no flag day proposals, no reference implementation in Bitcoin Core - Most Bitcoin Core developers treat quantum resistance as "important but not urgent" - The loudest debates are about property rights (what to do with exposed coins), not algorithms

5. Other Cryptocurrencies: Lessons for Bitcoin

5.1 Ethereum

Vitalik Buterin has been vocal about quantum preparedness. Ethereum's roadmap includes: - Account abstraction (EIP-4337) enabling wallets to swap signature schemes without protocol changes - Research into STARK-based signature verification (hash-based, quantum-resistant) - Ethereum's faster upgrade cadence means it can likely deploy PQC faster than Bitcoin

5.2 Algorand

Already uses Falcon-based signatures in its state proof protocol. Closest to a production PQC deployment among major chains.

5.3 QRL (Quantum Resistant Ledger)

Purpose-built quantum-resistant chain using XMSS (hash-based signatures). Small ecosystem but proves the concept works. Demonstrates the performance tradeoffs: larger transactions, slower verification.

5.4 Solana

No public quantum resistance roadmap. Solana's Ed25519 signatures have the same Shor vulnerability as Bitcoin's ECDSA. Higher transaction throughput means larger absolute impact from PQC signature bloat.

5.5 Lessons for Bitcoin

1. **Account abstraction helps** — Ethereum's ability to swap signature schemes per-account is a significant architectural advantage Bitcoin lacks
2. **Hash-based signatures work but are expensive** — QRL proves feasibility, also proves the performance cost is real
3. **First movers bear the risk** — Algorand chose FALCON early; if FALCON is broken, they're exposed. Bitcoin's conservatism may pay off.

6. Assessment and Recommendations

6.1 How Urgent Is This Really?

Medium urgency, high importance. The threat is not imminent (no CRQC exists or is close), but the remediation timeline is long (2–4+ years for a Bitcoin consensus change). The window for comfortable preparation is approximately **2026–2030**. After 2030, the timeline compresses uncomfortably.

The real risk is not a sudden quantum computer appearing overnight — it's the **gradual erosion of confidence**. As quantum hardware improves, even pre-CRQC milestones (breaking smaller elliptic curves, demonstrating Shor's on toy problems at scale) could trigger market panic and a rush to migrate that the network isn't prepared for.

6.2 What Should a Bitcoin Holder Do Today?

Action	Urgency	Difficulty
Never reuse addresses	Do now	Easy — most modern wallets do this by default
Use Taproot (P2TR) addresses	Do now	Easy — public key only exposed at spend time
Move coins from P2PK addresses	Soon	Medium — requires access to old keys
Avoid leaving large amounts in exposed addresses	Soon	Medium
Monitor BIP 360 progress	Ongoing	None — just stay informed
Consider cold storage for long-term holdings	Ongoing	Easy

Do NOT: Panic, sell, or move to altchains claiming quantum resistance. The timeline is measured in years, not months.

6.3 Most Likely Path Forward

1. **2026–2027:** BIP 360 undergoes technical review. Reference implementation developed. Community debate intensifies around the exposed-coins question.
2. **2027–2028:** Testnet deployment. Wallet developers begin adding P2QRH/P2MR support. Lightning Network evaluates PQC for channel establishment.
3. **2028–2029:** Activation debate. Miner signaling. Voluntary migration begins for early adopters.
4. **2029–2030:** Activation. New quantum-resistant address type available. Migration campaigns.
5. **2030+:** Community reckons with the exposed-coins question based on quantum hardware progress. Possible sunset rules for vulnerable address types.

This timeline assumes no quantum surprise. A credible demonstration of Shor's algorithm breaking a production-scale elliptic curve would compress everything dramatically.

Key Takeaways

1. **The threat is to signatures (ECDSA), not mining (SHA-256).** Quantum computing enables theft, not faster mining.
2. **~6.5 million BTC is already exposed.** These coins' public keys are visible on-chain and attackable offline whenever a CRQC arrives.
3. **FALCON-1024 is the leading candidate**, but signature sizes are 18× larger than ECDSA, creating real throughput concerns.
4. **BIP 360 exists and was merged in February 2026**, but activation is years away.
5. **The hardest problem is political, not technical**: what to do with Satoshi's coins and other exposed addresses.
6. **The preparation window is approximately 2026–2030.** After that, timelines compress uncomfortably.
7. **For individual holders: don't reuse addresses, use Taproot, and stay informed.** There's no need to panic.

References and Further Reading

- **BIP 360 (QuBit):** github.com/bitcoin/bips (merged February 2026)
- **Presidio Bitcoin Quantum Summit:** July 2025, 80 experts
- **Human Rights Foundation:** Quantum vulnerability exposure analysis (~6.51M BTC)
- **a16z Crypto:** "Quantum Computing and Cryptocurrency" (December 2025)
- **NIST PQC Standards:** FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA), FIPS 206 (FN-DSA)
- **ScienceDirect Survey:** "Navigating the quantum computing threat landscape for blockchains" (November 2025)

- **Google Willow:** 105-qubit chip with error correction milestone (2025)
- **BTQ Technologies:** ML-DSA Bitcoin node implementation (October 2025)
- **“The Pinnacle Architecture”:** Preprint suggesting sub-100K physical qubit ECDSA break (2026, unverified)

Changelog

Date	Update
2026-03-05	Initial report

This report is intended for informational purposes only and does not constitute investment or security advice. The field is evolving rapidly; verify claims against primary sources.